

PHASE RETRIEVAL FOR IMAGING PROBLEMS.

FAJWEL FOGEL, IRÈNE WALDSPURGER, AND ALEXANDRE D'ASPREMONT

ABSTRACT. We study convex relaxation algorithms for phase retrieval on imaging problems. We show that exploiting structural assumptions on the signal and the observations, such as sparsity, smoothness or positivity, can significantly speed-up convergence and improve recovery performance. We detail numerical results in molecular imaging experiments simulated using data from the Protein Data Bank (PDB).

1. INTRODUCTION

Phase retrieval seeks to reconstruct a complex signal, given a number of observations on the *magnitude* of linear measurements, i.e. solve

$$\begin{aligned} \text{find } & x \\ \text{such that } & |Ax| = b \end{aligned} \tag{1}$$

in the variable $x \in \mathbb{C}^p$, where $A \in \mathbb{R}^{n \times p}$ and $b \in \mathbb{R}^n$. This problem has direct applications in X-ray and crystallography imaging, diffraction imaging, Fourier optics or microscopy for example, in problems where physical limitations mean detectors usually capture the intensity of observations but cannot recover their phase. In what follows, we will focus on problems arising in diffraction imaging, where A is usually a Fourier transform, often composed with one or multiple masks (a technique sometimes called ptychography). The Fourier structure, through the FFT, considerably speeds up basic linear operations, which allows us to solve large scale convex relaxations on realistically large imaging problems. We will also observe that in many of the imaging problems we consider, the Fourier transform is very sparse, with *known support* (we lose the phase but observe the magnitude of Fourier coefficients), which allows us to considerably reduce the size of our convex phase retrieval relaxations.

Because the phase constraint $|Ax| = b$ is nonconvex, the phase recovery problem (1) is non-convex. Several greedy algorithms have been developed (see [Gerchberg and Saxton, 1972; Fienup, 1982; Griffin and Lim, 1984; Bauschke et al., 2002] among others), which alternate projections on the range of A and on the nonconvex set of vectors y such that $|y| = |Ax|$. While empirical performance is often good, these algorithms can stall in local minima. A convex relaxation was introduced in [Chai et al., 2011] and [Candes et al., 2011] (who call it PhaseLift) by observing that $|Ax|^2$ is a linear function of $X = xx^*$ which is a rank one Hermitian matrix, using the classical lifting argument for nonconvex quadratic programs developed in [Shor, 1987; Lovász and Schrijver, 1991]. The recovery of x is thus expressed as a rank minimization problem over positive semidefinite Hermitian matrices X satisfying some linear conditions, i.e. a matrix completion problem. This last problem has received a significant amount of attention because of its link to compressed sensing and the NETFLIX collaborative filtering problem. This minimum rank matrix completion problem is approximated by a semidefinite program which has been shown to recover x for several (random) classes of observation operators A [Candes et al., 2011, 2013a,b].

On the algorithmic side, [WalDSPURGER et al., 2012] showed that the phase retrieval problem (1) can be reformulated in terms of a single phase variable, which can be read as an extension of the MAXCUT combinatorial graph partitioning problem over the unit complex torus, allowing fast algorithms designed for solving semidefinite relaxations of MAXCUT to be applied to the phase retrieval problem.

Date: April 9, 2014.

2010 Mathematics Subject Classification. 94A12, 90C22, 90C27.

Key words and phrases. Phase recovery, semidefinite programming, X-ray diffraction, molecular imaging, Fourier optics.

On the experimental side, phase recovery is a classical problem in Fourier optics for example [Goodman, 2008], where a diffraction medium takes the place of a lens. This has a direct applications in X-ray and crystallography imaging, diffraction imaging or microscopy [Harrison, 1993; Bunk et al., 2007; Johnson et al., 2008; Miao et al., 2008; Dierolf et al., 2010].

Here, we implement and study several efficient convex relaxation algorithms for phase retrieval on imaging problem instances where A is based on a Fourier operator. We show in particular how structural assumptions on the signal and the observations (e.g. sparsity, smoothness, positivity, known support, oversampling, etc.) can be exploited to both speed-up convergence and improve recovery performance. While no experimental data is available from diffraction imaging problems with multiple randomly coded illuminations, we simulate numerical experiments of this type using molecular density information from the protein data bank [Berman et al., 2002]. Our results show in particular that the convex relaxation is stable and that in some settings, as few as two random illuminations suffice to reconstruct the image.

The paper is organized as follows. Section 2 briefly recalls the structure of some key algorithms used in phase retrieval. Section 3 describes applications to imaging problems and how structural assumptions can significantly reduce the cost of solving large-scale instances and improve recovery performance. Section 4 details some numerical experiments while Section 5 describes the interface to the numerical library developed for these problems.

Notations. We write \mathbf{S}_p (resp. \mathbf{H}_p) the cone of symmetric (resp. Hermitian) matrices of dimension p ; \mathbf{S}_p^+ (resp. \mathbf{H}_p^+) denotes the set of positive symmetric (resp. Hermitian) matrices. We write $\|\cdot\|_p$ the Schatten p -norm of a matrix, that is the p -norm of the vector of its eigenvalues (in particular, $\|\cdot\|_\infty$ is the spectral norm). We write A^\dagger the (Moore-Penrose) pseudoinverse of a matrix A , and $A \circ B$ the Hadamard (or componentwise) product of the matrices A and B . For $x \in \mathbb{R}^p$, we write $\mathbf{diag}(x)$ the matrix with diagonal x . When $X \in \mathbf{H}_p$ however, $\mathbf{diag}(X)$ is the vector containing the diagonal elements of X . For $X \in \mathbf{H}_p$, X^* is the Hermitian transpose of X , with $X^* = (\bar{X})^T$. Finally, we write b^2 the vector with components $b_i^2, i = 1, \dots, n$.

2. ALGORITHMS

In this section, we briefly recall several basic algorithmic approaches to solve the phase retrieval problem (1). Early methods were all based on extensions of an alternating projection algorithm. However, recent results showed that phase retrieval could be interpreted as a matrix completion problem similar to the NETFLIX problem, a formulation which yields both efficient convex relaxations and recovery guarantees.

2.1. Greedy algorithms. The phase retrieval problem (1) can be rewritten

$$\begin{aligned} & \text{minimize} && \|Ax - y\|_2^2 \\ & \text{subject to} && |y| = b \end{aligned} \tag{2}$$

where we now optimize over both phased observations $y \in \mathbb{C}^n$ and signal $x \in \mathbb{C}^p$. Several greedy algorithms attempt to solve this problem using variants of alternating projections, one iteration minimizing the quadratic error (the objective of (2)), the next normalizing the moduli (to satisfy the constraint). We detail some of the most classical examples in the paragraphs that follow.

The algorithm **Gerchberg-Saxton** by [Gerchberg and Saxton, 1972] for instance seeks to reconstruct $y = Ax$ and alternates between orthogonal projections on the range of A and normalization of the magnitudes $|y|$ to match the observations b . The cost per iteration of this method is minimal but convergence (when it happens) is often slow.

A classical “input-output” variant, detailed here as algorithm **Fienup**, introduced by [Fienup, 1982] adds an extra penalization step which usually speeds up convergence and improves recovery performance when additional information is available on the support of the signal. Oversampling the Fourier transform forming

Algorithm 1 Gerchberg-Saxton.

Input: An initial $y^1 \in \mathbf{F}$, i.e. such that $|y^1| = b$.

1: **for** $k = 1, \dots, N - 1$ **do**

2: Set

$$y_i^{k+1} = b_i \frac{(AA^\dagger y^k)_i}{|(AA^\dagger y^k)_i|}, \quad i = 1, \dots, n. \quad (\text{Gerchberg-Saxton})$$

3: **end for**

Output: $y_N \in \mathbf{F}$.

Algorithm 2 Fienup

Input: An initial $y^1 \in \mathbf{F}$, i.e. such that $|y^1| = b$, a parameter $\beta > 0$.

1: **for** $k = 1, \dots, N - 1$ **do**

2: Set

$$w_i = \frac{(AA^\dagger y^k)_i}{|(AA^\dagger y^k)_i|}, \quad i = 1, \dots, n.$$

3: Set

$$y_i^{k+1} = y_i^k - \beta(y_i^k - b_i w_i) \quad (\text{Fienup})$$

4: **end for**

Output: $y_N \in \mathbf{F}$.

A in imaging problems usually helps performance as well. Of course, in all these cases, convergence to a global optimum cannot be guaranteed but empirical recovery performance is often quite good.

2.2. PhaseLift: semidefinite relaxation in signal. Using a classical lifting argument by [Shor, 1987], and writing

$$|a_i^* x|^2 = b_i^2 \iff \text{Tr}(a_i a_i^* x x^*) = b_i^2$$

[Chai et al., 2011; Candes et al., 2011] reformulate the phase recovery problem (1) as a matrix completion problem, written

$$\begin{aligned} & \text{minimize} && \mathbf{Rank}(X) \\ & \text{subject to} && \text{Tr}(a_i a_i^* X) = b_i^2, \quad i = 1, \dots, n \\ & && X \succeq 0 \end{aligned}$$

in the variable $X \in \mathbf{H}_p$, where $X = x x^*$ when exact recovery occurs. This last problem can be relaxed as

$$\begin{aligned} & \text{minimize} && \text{Tr}(X) \\ & \text{subject to} && \text{Tr}(a_i a_i^* X) = b_i^2, \quad i = 1, \dots, n \\ & && X \succeq 0 \end{aligned} \quad (\text{PhaseLift})$$

which is a semidefinite program (called **PhaseLift** by Candes et al. [2011]) in the variable $X \in \mathbf{H}_p$. This problem is solved in [Candes et al., 2011] using first order algorithms implemented in [Becker et al., 2011]. This semidefinite relaxation has been shown to recover the true signal x exactly for several classes of observation operators A [Candes et al., 2011, 2013a,b].

2.3. PhaseCut: semidefinite relaxation in phase. As in [Waldspurger et al., 2012] we can rewrite the phase reconstruction problem (1) in terms of a phase variable u (such that $|u| = 1$) instead of the signal x . In the noiseless case, we then write the constraint $|Ax| = b$ as $Ax = \text{diag}(b)u$, where $u \in \mathbb{C}^n$ is a phase vector, satisfying $|u_i| = 1$ for $i = 1, \dots, n$, so problem (1) becomes

$$\begin{aligned} & \text{minimize} && \|Ax - \text{diag}(b)u\|_2^2 \\ & \text{subject to} && |u_i| = 1 \end{aligned} \quad (3)$$

where we optimize over both phase $u \in \mathbb{C}^n$ and signal $x \in \mathbb{C}^p$. While the objective of this last problem is jointly convex in (x, u) , the phase constraint $|u_i| = 1$ is not.

Now, given the phase, signal reconstruction is a simple least squares problem, i.e. given u we obtain x as

$$x = A^\dagger \mathbf{diag}(b)u \quad (4)$$

where A^\dagger is the pseudo inverse of A . Replacing x by its value in (3), the phase recovery problem becomes

$$\begin{aligned} &\text{minimize} && u^* M u \\ &\text{subject to} && |u_i| = 1, \quad i = 1, \dots, n, \end{aligned} \quad (5)$$

in the variable $u \in \mathbb{C}^n$, where the Hermitian matrix

$$M = \mathbf{diag}(b)(\mathbf{I} - AA^\dagger) \mathbf{diag}(b)$$

is positive semidefinite. This problem is non-convex in the phase variable u . [Waldspurger et al., 2012] detailed greedy algorithm **Greedy** to locally optimize (5) in the phase variable.

Algorithm 3 Greedy algorithm in phase.

Input: An initial $u \in \mathbb{C}^n$ such that $|u_i| = 1, i = 1, \dots, n$. An integer $N > 1$.

- 1: **for** $k = 1, \dots, N$ **do**
- 2: **for** $i = 1, \dots, n$ **do**
- 3: Set

$$u_i = \frac{-\sum_{j \neq i} M_{ji} \bar{u}_j}{\left| \sum_{j \neq i} M_{ji} \bar{u}_j \right|} \quad (\text{Greedy})$$

- 4: **end for**
- 5: **end for**

Output: $u \in \mathbb{C}^n$ such that $|u_i| = 1, i = 1, \dots, n$.

A convex relaxation to (5) was also derived in [Waldspurger et al., 2012] using the classical lifting argument for nonconvex quadratic programs developed in [Shor, 1987; Lovász and Schrijver, 1991]. This relaxation is written

$$\begin{aligned} &\text{min.} && \mathbf{Tr}(UM) \\ &\text{subject to} && \mathbf{diag}(U) = \mathbf{1}, U \succeq 0, \end{aligned} \quad (\text{PhaseCut})$$

which is a semidefinite program (SDP) in the matrix $U \in \mathbf{H}_n$. This problem has a structure similar to the classical MAXCUT relaxation and instances of reasonable size can be solved using specialized implementations of interior point methods designed for that problem [Helmberg et al., 1996]. Larger instances are solved in [Waldspurger et al., 2012] using the block-coordinate descent algorithm **BlockPhaseCut**.

Ultimately, algorithmic choices heavily depend on problem structure, and these will be discussed in detail in the section that follows. In particular, we will study how to exploit structural information on the signal (nonnegativity, sparse 2D FFT, etc.), to solve realistically large instances formed in diffraction imaging applications.

3. IMAGING PROBLEMS

In the imaging problems we study here, various illuminations of a single object are performed through randomly coded masks, hence the matrix A is usually formed using a combination of random masks and Fourier transforms, and we have significant structural information on both the signal we seek to reconstruct (regularity, etc.) and the observations (power law decay in frequency domain, etc.). Many of these additional structural hints can be used to speedup numerical operations, convergence and improve phase retrieval performance. The paragraphs that follow explore these points in more detail.

Algorithm 4 Block Coordinate Descent Algorithm for **PhaseCut**.

Input: An initial $U^0 = \mathbf{I}_n$ and $\nu > 0$ (typically small). An integer $N > 1$.

- 1: **for** $k = 1, \dots, N$ **do**
- 2: Pick $i \in [1, n]$.
- 3: Compute

$$u = U_{i^c, i}^k M_{i^c, i} \quad \text{and} \quad \gamma = u^* M_{i^c, i} \quad (\text{BlockPhaseCut})$$

- 4: If $\gamma > 0$, set

$$U_{i^c, i}^{k+1} = U_{i, i^c}^{k+1*} = -\sqrt{\frac{1-\nu}{\gamma}} x$$

else

$$U_{i^c, i}^{k+1} = U_{i, i^c}^{k+1*} = 0.$$

- 5: **end for**

Output: A matrix $U \succeq 0$ with $\text{diag}(U) = 1$.

3.1. Fourier operators. In practical applications, because of the structure of the linear operator A , we may often reduce numerical complexity, using the Fourier structure of A to speedup the single matrix-vector product in algorithm **BlockPhaseCut**. We detail the case where A corresponds to a Fourier transform combined with k random masks, writing $I_1, \dots, I_k \in \mathbb{C}^p$ be the illumination masks. The image by A of some signal $x \in \mathbb{C}^p$ is then

$$Ax = \begin{pmatrix} \mathcal{F}(I_1 \circ x) \\ \vdots \\ \mathcal{F}(I_k \circ x) \end{pmatrix},$$

and the pseudo-inverse of A also has a simple structure, with

$$A^\dagger \begin{pmatrix} y_1 \\ \vdots \\ y_k \end{pmatrix} = \sum_{l=1}^k \mathcal{F}^{-1}(y_l) \circ I'_l$$

where I'_l is the dual filter of I_l , which is

$$I'_l = \bar{I}_l / \left(\sum_s |I_s|^2 \right).$$

With the fast Fourier transform, computing the image of a vector by A or A^\dagger only requires $O(kp \log(p))$ floating-point operations. For any $v \in \mathbb{C}^n$, $Mv = \text{diag}(b)(\mathbf{I} - AA^\dagger) \text{diag}(b)v$ may then be computed using $O(kp \log p)$ operations instead of $O(k^2 p^2)$ for naive matrix-vector multiplications.

In algorithms **Greedy** and **BlockPhaseCut**, we also need to extract quickly columns from M without having to store the whole matrix. Extracting the column corresponding to index i in block $l \leq k$ reduces to the computation of $AA^\dagger \delta_{i,l}$ where $\delta_{i,l} \in \mathbb{C}^{kp}$ is the vector whose coordinates are all zero, except the i -th one of l -th block. If we write $\delta_i \in \mathbb{C}^p$ the Dirac in i , the preceding formulas yields

$$AA^\dagger \delta_{i,l} = \begin{pmatrix} \delta_i \star \mathcal{F}(I_1 \circ I'_l) \\ \vdots \\ \delta_i \star \mathcal{F}(I_k \circ I'_l) \end{pmatrix}.$$

Convolution with δ_i is only a shift and vectors $\mathcal{F}(I_s \circ I'_l)$ may be precomputed so this operation is very fast.

3.2. Low rank iterates. In instances where exact recovery occurs, the solution to the semidefinite programming relaxation (**PhaseCut**) has rank one. It is also likely to have low rank in a neighborhood of the optimum. This means that we can often store a compressed version of the iterates U in algorithm **Block-PhaseCut** in the form of their low rank approximation $U = VV^*$ where $V \in \mathbb{C}^{n \times k}$. Each iteration updates a single row/column of U which corresponds to a rank two update of U , hence updating the SVD means computing a few leading eigenvalues of the matrix $VV^* + LL^*$ where $L \in \mathbb{C}^{n \times 2}$. This update can be performed using Lanczos type algorithms and has complexity $O(kn \log n)$. Compressed storage of U saves memory and also speeds-up the evaluation of the vector matrix product $U_{i^c, i^c} M_{i^c, i}$ which costs $O(nk)$ given a decomposition $U_{i^c, i^c} = VV^*$, instead of $O(n^2)$ using a generic representation of the matrix U .

3.3. Bounded support. In many inverse problems the signal we are seeking to reconstruct is known to be sparse in some basis and exploiting this structural information explicitly usually improves signal recovery performance. This is for example the basis of compressed sensing where ℓ_1 penalties encourage sparsity and provide recovery guarantees when the true signal is actually sparse.

The situation is a lot simpler in some of the molecular imaging problems we are studying below since the electron density we are trying to recover is often smooth, which means that its Fourier transform will be sparse, *with known support*. While we lose the phase, we do observe the magnitude of the Fourier coefficients so we can rank them by magnitude. This allows us to considerably reduce the size of the SDP relaxation without losing much reconstruction fidelity, i.e. in many cases we observe that a significant fraction of the coefficients of b are close to zero. From a computational point of view, sparsity in b allows us to solve a truncated semidefinite relaxation (**PhaseCut**). See Figure 1 for an illustration of this phenomenon on the caffeine molecule.




FIGURE 1. Electronic density for the caffeine molecule (left), its 2D FFT transform (diffraction pattern, center), the density reconstructed using 2% of the coefficients with largest magnitude in the FFT (right).

Indeed, without loss of generality, we can reorder the observations b such that we approximately have $b = (b_1^T, 0)^T$. Similarly, we note

$$u = \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}, \quad A = \begin{pmatrix} A_1 \\ A_2 \end{pmatrix}, \quad A^\dagger = ((A^\dagger)_1 \quad (A^\dagger)_2).$$

Using the fact that $b_2 = 0$, the matrix M in the objective of (5) can itself be written in blocks, that is

$$M = \begin{pmatrix} M_1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Since $b_2 = 0$, any complex vector with coefficients of magnitude one can be taken for u_2 and the optimization problem (5) is equivalent to

$$\begin{aligned} & \text{minimize} && u_1^* M_1 u_1 \\ & \text{subject to} && |u_{1_i}| = 1, \quad i = 1, \dots, n, \end{aligned} \tag{6}$$

in the variable $u_1 \in \mathbb{C}^{n_1}$, where the Hermitian matrix

$$M_1 = \mathbf{diag}(b_1)(\mathbf{I} - A_1(A_1^\dagger)_1) \mathbf{diag}(b_1)$$

is positive semidefinite. This problem can in turn be relaxed into a PhaseCut problem which is usually considerably smaller than the original (PhaseCut) problem since M_1 is typically a fraction of the size of M .

3.4. Real, positive densities. In some cases, such as imaging experiments where a random binary mask is projected on an object for example, we know that the linear observations are formed as the Fourier transform of a *positive* measure. This introduces additional restrictions on the structure of these observations, which can be written as convex constraints on the phase vector. We detail two different ways of accounting for this positivity assumption.

3.4.1. Direct nonnegativity constraints on the density. In the case where the signal is real and nonnegative, [Waldspurger et al., 2012] show that problem (3) can be modified to specifically account for the fact that the signal is real, by writing it

$$\min_{\substack{u \in \mathbb{C}^n, |u_i|=1, \\ x \in \mathbb{R}^p}} \|Ax - \mathbf{diag}(b)u\|_2^2,$$

using the operator $\mathcal{T}(\cdot)$ defined as

$$\mathcal{T}(Z) = \begin{pmatrix} \operatorname{Re}(Z) & -\operatorname{Im}(Z) \\ \operatorname{Im}(Z) & \operatorname{Re}(Z) \end{pmatrix} \quad (7)$$

we can rewrite the phase problem on real valued signal as

$$\begin{aligned} & \text{minimize} && \left\| \mathcal{T}(A) \begin{pmatrix} x \\ 0 \end{pmatrix} - \mathbf{diag} \begin{pmatrix} b \\ b \end{pmatrix} \begin{pmatrix} \operatorname{Re}(u) \\ \operatorname{Im}(u) \end{pmatrix} \right\|_2^2 \\ & \text{subject to} && u \in \mathbb{C}^n, |u_i| = 1 \\ & && x \in \mathbb{R}^p. \end{aligned}$$

The optimal solution of the inner minimization problem in x is given by $x = A_2^\dagger B_2 v$, where

$$A_2 = \begin{pmatrix} \operatorname{Re}(A) \\ \operatorname{Im}(A) \end{pmatrix}, \quad B_2 = \mathbf{diag} \begin{pmatrix} b \\ b \end{pmatrix}, \quad \text{and} \quad v = \begin{pmatrix} \operatorname{Re}(u) \\ \operatorname{Im}(u) \end{pmatrix}$$

hence the problem is finally rewritten

$$\begin{aligned} & \text{minimize} && \|(A_2 A_2^\dagger B_2 - B_2)v\|_2^2 \\ & \text{subject to} && v_i^2 + v_{n+i}^2 = 1, \quad i = 1, \dots, n, \end{aligned}$$

in the variable $v \in \mathbb{R}^{2n}$. This can be relaxed as above by the following problem

$$\begin{aligned} & \text{minimize} && \mathbf{Tr}(VM_2) \\ & \text{subject to} && V_{ii} + V_{n+i, n+i} = 1, \quad i = 1, \dots, n, \\ & && V \succeq 0, \end{aligned} \quad (\text{PhaseCutR})$$

which is a semidefinite program in the variable $V \in \mathbf{S}_{2n}$, where

$$M_2 = (A_2 A_2^\dagger B_2 - B_2)^T (A_2 A_2^\dagger B_2 - B_2) = B_2^T (\mathbf{I} - A_2 A_2^\dagger) B_2.$$

Because $x = A_2^\dagger B_2 v$ for real instances, we can add a nonnegativity constraint to this relaxation, using

$$xx^T = (A_2^\dagger B_2)uu^T(A_2^\dagger B_2)^T$$

and the relaxation becomes

$$\begin{aligned} & \text{minimize} && \mathbf{Tr}(VM_2) \\ & \text{subject to} && (A_2^\dagger B_2)V(A_2^\dagger B_2)^T \succeq 0, \\ & && V_{ii} + V_{n+i, n+i} = 1, \quad i = 1, \dots, n, \\ & && V \succeq 0, \end{aligned}$$

which is a semidefinite program in $V \in \mathbf{S}_{2n}$.

3.4.2. *Bochner's theorem and the Fourier transform of positive measures.* Another way to include nonnegativity constraints on the signal, which preserves some of the problem structure, is to use Bochner's theorem. Recall that a function $f : \mathbb{R}^s \mapsto \mathbb{C}$ is *positive semidefinite* if and only if the matrix B with coefficients $B_{ij} = f(x_i - x_j)$ is Hermitian positive semidefinite for any sequence $x_i \in \mathbb{R}^s$. Bochner's theorem then characterizes Fourier transforms of positive measures.

Theorem 3.1. (Bochner) *A function $f : \mathbb{R}^s \mapsto \mathbb{C}$ is positive semidefinite if and only if it is the Fourier transform of a (finite) nonnegative Borel measure.*

Proof. See [Berg et al., 1984] for example. ■

For simplicity, we first illustrate this in dimension one. Suppose that we observe the magnitude of the Fourier transform of a discrete nonnegative signal $x \in \mathbb{R}^p$ so that

$$|\mathcal{F}x| = b$$

with $b \in \mathbb{R}^n$. Our objective now is to reconstruct a phase vector $u \in \mathbb{C}^n$ such that $|u| = 1$ and

$$\mathcal{F}x = \mathbf{diag}(b)u.$$

If we define the Toeplitz matrix

$$B_{ij}(y) = y_{|i-j|+1}, \quad 1 \leq j \leq i \leq p,$$

so that

$$B(y) = \begin{pmatrix} y_1 & y_2^* & \cdots & \cdots & y_n^* \\ y_2 & y_1 & y_2^* & \cdots & \vdots \\ \vdots & y_2 & y_1 & y_2^* & \vdots \\ \vdots & \vdots & \ddots & \ddots & \ddots \\ y_n & \cdots & y_2 & y_1 & y_2^* \end{pmatrix}$$

then when $\mathcal{F}x = \mathbf{diag}(b)u$, Bochner's theorem states that $B(\mathbf{diag}(b)u) \succeq 0$ iff $x \geq 0$. The constraint $B(\mathbf{diag}(b)u) \succeq 0$ is a linear matrix inequality in u , hence is convex.

Suppose that we observe multiple illuminations and that the k masks $I_1, \dots, I_k \in \mathbb{R}^{p \times p}$ are also nonnegative (e.g. random coherent illuminations), we have

$$Ax = \begin{pmatrix} \mathcal{F}(I_1 \circ x) \\ \vdots \\ \mathcal{F}(I_k \circ x) \end{pmatrix},$$

and the phase retrieval problem (5) for positive signals x is now written

$$\begin{aligned} & \text{minimize} && u^* M u \\ & \text{subject to} && B_j(\mathbf{diag}(b)u) \succeq 0, \quad j = 1, \dots, k \\ & && |u_i| = 1, \quad i = 1, \dots, n, \end{aligned}$$

where $B_j(y)$ is the matrix $B(y^{(j)})$, where $y^{(j)} \in \mathbb{C}^p$ is the j^{th} subvector of y (one for each of the k masks). We can then adapt the **PhaseCut** relaxation to incorporate the positivity requirement. In the one dimensional case, using again the classical lifting argument in [Shor, 1987; Lovász and Schrijver, 1991], it becomes

$$\begin{aligned} & \text{min.} && \mathbf{Tr}(UM) \\ & \text{subject to} && \mathbf{diag}(U) = 1, \quad u_1 = 1, \\ & && B_j(\mathbf{diag}(b)u) \succeq 0, \quad j = 1, \dots, k \\ & && \begin{pmatrix} U & u \\ u^* & 1 \end{pmatrix} \succeq 0 \end{aligned} \tag{PhaseCut+}$$

	Caffeine	Cocaine	Lysozyme
Nb. atoms	14	43	1309
16 × 16 res.	58 %	40 %	11 %
32 × 32 res.	44 %	40 %	20 %
64 × 64 res.	15 %	55 %	14 %
128 × 128 res.	4 %	55 %	4 %

TABLE 1. Percentage of 2D FFT coefficients required to reach $10^{-1.5}$ relative MSE, without oversampling.

in the variables $U \in \mathbf{S}_n$ and $u \in \mathbb{C}^n$. The phase vector u is fixed up to an arbitrary global shift, and the additional constraint $u_1 = 1$ allows us to exclude degenerate solutions with $u = 0$. Similar results apply in multiple dimensions, since the 2D Fourier transform is simply computed by applying the 1D Fourier transform first to columns then to rows.

The SDP relaxation **PhaseCut+** cannot be solved using block coordinate descent. Without positivity constraints, the relaxation **PhaseCutR** designed for real signals can be solved efficiently using the algorithm in [Helmberg et al., 1996]. The constraint structure in **PhaseCutR** means that the most expensive step at each iteration of the algorithm in [Helmberg et al., 1996] is computing the inverse of a symmetric matrix of dimension n (or less, exploiting sparsity in b). Sparse instances of the more complex relation **PhaseCut+** were solved using SDPT3 [Toh et al., 1999] in what follows.

4. NUMERICAL EXPERIMENTS

We study molecular imaging problems based on electronic densities obtained from the Protein Data Bank [Berman et al., 2002]. From a 3D image, we obtain a 2D projection by integrating the third dimension. After normalizing these images, we simulate multiple diffraction observations for each molecule, using several random masks. Here, our masks consist of randomly generated binary filters placed before the sample, but other settings are possible [Candes et al., 2013b]. Our vector of observations then corresponds to the magnitude of the Fourier transform of the componentwise product of the image and the filter. As in the SPSIM package [Maia, 2013] simulating diffraction imaging experiments, random Poisson noise is added to the observations, modeling sensor and electronic noise. More specifically, the noisy intensity measurements are obtained using the following formula,

$$I = \sqrt{\max \left\{ 0, \alpha \cdot \text{Poisson} \left(\frac{|Ax|^2}{\alpha} \right) \right\}},$$

where α is the input level of noise, and $\text{Poisson}(\lambda)$ is a random Poisson sample of mean λ . We ensure that all points of the electronic density are illuminated at least once by the random masks (the first mask lets all the signal go through) and call mask “resolution” the number of pixels in a square unit of the mask. For instance masks of resolution 4×4 pixels in a 16×16 pixels image will consist of sixteen square blocks of size 4×4 pixels, each block being either all zeros or all ones.

We present numerical experiments on two molecules from the Protein Data Bank (PDB), namely caffeine and cocaine, with very different structure (properly projected, caffeine is mostly circular, cocaine has a star shape). Images of the caffeine and cocaine molecules at low and high resolutions are presented in Figure 2. We first use “high” 128×128 pixels resolutions to evaluate the sensitivity of **PhaseCut** to noise and number of masks using the fast **BlockPhaseCut** algorithm (see section 4.1). We then use a “low” 16×16 pixels image resolution to compare **PhaseCut** formulations using structural constraints, *i.e.* complex **PhaseCut**, real **PhaseCutR**, and **PhaseCut+** (with positivity constraints, see section 4.2) on a large number of random experiments.




FIGURE 2. Two molecules, caffeine (left) and cocaine (right), at two resolutions: 16×16 and 128×128 .

	Caffeine	Cocaine	Lysozyme
Nb. atoms	14	43	1309
16×16 res.	48 %	34 %	10 %
32×32 res.	37 %	35 %	17 %
64×64 res.	13 %	48 %	12 %
128×128 res.	4 %	49 %	4 %

TABLE 2. Percentage of 2D FFT coefficients required to reach $10^{-1.5}$ relative MSE, with 2x oversampling.

4.1. High resolution experiments using BlockPhaseCut. We first compare the results obtained by the **Fienup** and **BlockPhaseCut** algorithms while varying the number of masks and the noise level. For the PhaseCut relaxation, in order to deal with the large size of the lifted matrix, we use the low rank approximation described in §3.2 to store iterates and exploit sparsity in the magnitude of the observations vector described as described in §3.3. Tables 1 and 2 illustrate the impact of image resolution and oversampling on the fraction of coefficients required to approximately reconstruct a molecular density up to a given quality threshold, for various molecules. We observe that the sparsity of 2D FFTs increases with resolution and oversampling, but varies from one molecule to another. We then retrieve the phase vector as the first eigenvector in the final low rank approximation, then refine it with the greedy algorithms **Greedy** or **Fienup**.

4.1.1. Parameters. More specifically, in the experiments that follow, the image was of size 128×128 , we used a rank of two for the low rank approximation, kept the largest 1000 observations, did 5000 iterations of algorithm **Fienup**, and 20 cycles of algorithm **BlockPhaseCut** (one cycle corresponds to optimizing once over all rows/columns of the lifted matrix). We compared the results of the phase recovery using one to four masks, and three different levels of Poisson noise (no noise, “small” noise, “large” noise). In all settings, all points of the electronic density were illuminated at least once by the random masks (the first mask lets all the signal go through). The noisy (Poisson) intensity measurements were obtained using the formula described above. Experiments were performed on a regular Linux desktop using Matlab for the greedy algorithms and a C implementation of the block coordinate algorithm for PhaseCut. Reported CPU times are in seconds.

4.1.2. Results. In most cases both algorithm **Fienup** and **BlockPhaseCut** seem to converge to the (global) optimal solution, though **Fienup** is much faster. In some cases however, such as the experiment with two filters and no noise in Figure 3, initializing algorithm **Fienup** with the solution from **BlockPhaseCut** significantly outperforms the solution obtained by algorithm **Fienup** alone, which appears to be stuck in a local

minimum. The corresponding MSE values are listed in Table 3. In Figure 5 we plot the histogram of MSE for the noiseless case with only two illuminations, using either algorithm **Fienup**, or **BlockPhaseCut** followed by greedy refinements, over many random illumination configurations. We observe that in many samples, algorithm **Fienup** gets stuck in a local optimum, while the SDP always converges to a global optimum.




FIGURE 3. Solution of the semidefinite relaxation algorithm **BlockPhaseCut** followed by greedy refinements, for various values of the number of filters and noise level α .

4.2. Performance of PhaseCut relaxations with respect to number of masks, noise and filter resolution.

We now compare PhaseCut formulations with structural constraints, *i.e.* complex **PhaseCut**, real **PhaseCutR**, and **PhaseCut+** (with positivity constraints, see section 4.2) on a large number of random experiments formed using “low” 16×16 pixels image resolution.

4.2.1. *Varying the number of masks.* Masks are of resolution 1×1 and no noise is added. As shown in Figures 6 and 7, **PhaseCut**, **PhaseCutR** and **PhaseCut+** with **Fienup** post processing (respectively “SDP + Fienup HIO”, “SDP + real + Fienup HIO”, “SDP + real + toeplitz + Fienup HIO” and “Fienup HIO” curves on the figure) all outperform **Fienup** alone. For **PhaseCut**, in most cases, two to three masks seem enough to exactly recover the phase. Moreover, as expected, **PhaseCutR** performs a little bit better than **PhaseCut**, but surprisingly, positivity constraints of **PhaseCut+** do not seem to improve the solution of **PhaseCutR** in




FIGURE 4. Solution of the greedy algorithm **Fienup**, for various values of the number of filters and noise level α .




FIGURE 5. Histogram of MSE for the noiseless case with only two illuminations, using either algorithm **Fienup** (red), or **BlockPhaseCut** (blue) followed by greedy refinements, over many random starts.

these experiments. Finally, as shown in Figures 8 and 9, oversampling the Fourier transform seems to have a positive impact on the reconstruction. Results on caffeine and cocaine are very similar.

Nb masks	α	SDP obj	SDP Refined obj	Fienup obj	SDP time	Fienup time
1	0	0.410	0.021	0.046	104	78
2	0	2.157	0.069	0.877	179	156
3	0	4.036	0.000	0.000	256	241
4	0	7.075	0.000	0.000	343	343
1	10^{-3}	0.806	0.691	0.704	102	67
2	10^{-3}	3.300	2.986	2.992	182	144
3	10^{-3}	5.691	5.276	5.277	263	229
4	10^{-3}	8.074	7.259	7.259	351	308
1	10^{-2}	3.299	3.187	3.193	104	71
2	10^{-2}	7.492	8.622	8.646	182	145
3	10^{-2}	10.419	14.277	14.288	263	222
4	10^{-2}	14.139	21.472	21.444	349	305

TABLE 3. Performance comparison between algorithms **Fienup** and **BlockPhaseCut** for various values of the number of filters and noise level α . CPU times are in seconds.

4.2.2. *Varying mask resolution.* Here, two or three masks are used and no noise is added. As shown in Figures 10, 11, 12 and 13, we can see that the MSE of reconstructed images increase with the resolution of masks. Moreover **PhaseCutR** is more robust to lower mask resolution than **PhaseCut**. Finally, as expected, with more randomly masked illuminations, we can afford to lower mask resolution.

4.2.3. *Varying noise levels.* Here two masks are used (the minimum), with resolution 1×1 . Poisson noise is added (parameterized by α). As shown in Figures 14, and 15, we can see that **PhaseCut** and **PhaseCutR** are stable with regards to noise, i.e. we obtain a linear increase of the log MSE with respect to the log noise.

5. USER GUIDE

We provide here the instructions to artificially recover the image of a molecule from the Protein Data Bank using **PhaseCutToolbox** (download at www.di.ens.fr/~aspremon). This example is entirely reproduced with comments in the script `testPhaseCut.m`.

5.1. **Installation.** Our toolbox works on all recent versions of MATLAB on Mac OS X, and on MATLAB versions anterior to 2008 on Linux (there might be conflicts with Arpack library for ulterior versions, when using **BlockPhaseCut**). Installation only requires to put the toolbox folder and subdirectories on the Matlab path. Use for instance the command:

```
>> addpath(genpath('MYPATH/PhaseCutToolbox'));
```

where MYPATH is the directory where you have copied the toolbox.

5.2. **Generate the diffraction pattern of a molecule.** Suppose we work with the caffeine molecule, on an image of resolution 128×128 pixels. We set the corresponding input variables.

```
>> nameMol='caffeine.pdb';
>> N = 128 ;
```

Now, we set the parameters of the masks. The number of masks (also called filters or illuminations) is set to 2. Moreover we set the filter resolution to 1. The filter resolution corresponds to the square root of the

number of pixels in each block of the binary filter. The filter resolution must divide N (the square root of the number of pixels in the image).

```
>> filterRes = 1 ;
>> nb_filters=2;
```

Since the filters are generated randomly, we set the seed of the uniform random generator to 1 in order to get reproducible experiments. Note that the quality of the phase retrieval may depend on the shape of the generated masks, especially when using only 2 or 3 filters.

```
>> rand('seed',1);
```

Now we can generate an image, 2 masks and their corresponding diffraction patterns. We set the level of noise on the observations to zero here (*i.e.* no noise). α is the level of Poisson noise, and β is the level of Gaussian noise.

```
>> alpha=0;
>> beta=0;
```

We set the oversampling parameter for the Fourier transform to 2.

```
>> OSF = 2;
```

The total number of observations, *i.e.* the size of the vector b is

```
>> nbObs=N*N*OSF*OSF*nb_filters;
```

Suppose that we want to use only the first largest one thousand observations in PhaseCut, we set

```
>> nbObsKept=1000;
```

Note that the number of observations that is sufficient to get close to the optimal solution depends on the size of the data N and the sparsity of the vector b . From a more practical point of view, the larger `nbObsKept`, the more time intensive the optimization. Therefore, for a quick test we recommend setting `nbObsKept` to a few thousands, then increasing it if the results are not satisfying.

Finally we call the function `genData` which is going to generate both the image x we want to recover, `filters`, and observations b . `bs` corresponds to the thousand largest observations, `xs` is the image recovered with the true phase but using only `bs`. `idx_bs` is the logical indicator vector of `bs` (`bs=b(idx_bs)`). We put `displayFig` to 1 in order to display the filters, the images of the molecule x and xs , as well as the diffraction patterns (with and without noise).

```
>> displayFig=1;
>> [x,b,filters,bs,xs,idx_bs] = genData(nameMol, nb_filters, ...
    filterRes, N, alpha, beta, OSF, nbObsKept, displayFig);
```

5.3. Phase Retrieval using Fienup and/or PhaseCut. Using the data generated in the previous section, we retrieve the phase of the observations vector b . Suppose we want to use the SDP relaxation with greedy refinement, we set

```
>> method='SDPRefined';
```

The other choices for `method` are `'Fienup'`, `'Fienup HIO'` and `'SDP'` (no greedy refinement). We set the initial (full) phase vector u to the vector of ones, and the number of iterations for Fienup algorithm to 5000. The number of iterations for Fienup algorithm must be large enough so that the objective function converges to a stationary point. In most cases 5000 iterations seems to be enough.

```
>> param.uInit=ones(nbObs,1);
>> param.nbIterFienup=5000;
```

We also need to choose which algorithm we want to use in order to solve the SDP relaxation. For high resolution images, we recommend to always use the block coordinate descent algorithm with a low rank approximation of the lifted matrix (BCDLR), since interior points methods (when using SDPT3 or Mosek) and block coordinate descent without low rank approximation (BCD) become very slow when the number of observations used is over a few thousands.

```
>> param.SDPsolver='BCDLR';
```

If we had wanted to solve **PhaseCutR** or **PhaseCut+** we would have set

```
>> param.SDPsolver='realSDPT3';
```

or

```
>> param.SDPsolver='ToepSDPT3';
```

We can now set up the parameters for the BCDLR solver.

```
>> param.nbCycles=20;
```

```
>> param.r=2;
```

One cycle corresponds to optimizing over all the columns of the lifted matrix. In most cases, it seems that using `nbCycles` between 20 and 40 is enough to get close to the optimum, at least when refining the solution with Fienup algorithm. `r` is the rank for the low rank approximation of the lifted matrix. Similarly it seems that `r` between 2 and 4 gives reasonable results. Note that you can check that the low rank approximation is valid by looking at the maximum ratio between the last and the first eigenvalues throughout all iterations of the BCDLR algorithm. This ratio is outputted as `relax.eigRatio` when calling the function `retrievePhase` (see below). We finally call the function `retrievePhase` in order to solve the SDP relaxation with greedy refinement.

```
>> data.b=b;
```

```
>> data.bs=bs;
```

```
>> data.idx_bs=idx_bs;
```

```
>> data.OSF=OSF;
```

```
>> data.filters=filters;
```

```
>> [retrievedPhase, objValues, finalObj,relax] = retrievePhase(data,method,param);
```

The function `retrievePhase` outputs the vector of retrieved phase as `retrievedPhase` and the values of the objective function at each iteration/cycle of the algorithm in `objValues` (add `.Fienup`, `.SDP` `.SDPRefined` to `retrievedPhase` and `objValues` to get the corresponding retrieved phase and objective value). If using the SDP relaxation, the vector `retrievedPhase` is the first eigenvector of the final lifted matrix in `PhaseCut`. Note that the objective value in Fienup and in the SDP relaxation do not correspond exactly since the lifted matrix may be of rank bigger than one during the iterations of the BCDLR. Therefore we also output `finalObj`, which is the objective value of the phase vector extracted from the lifted matrix (i.e. the vector `retrievedPhase`). The image can now be retrieved using the command

```
>> xRetrieved=pseudo_inverse_A(retrievedPhase.SDPRefined.*b,filters,M);
```

Finally you can visualize the results using the following standard Matlab commands, plotting the objective values

```
>> figure(1)
```

```
>> subplot(2,1,1);
```

```
>> title(method)
```

```
>> plot(log10(abs(objValues))); axis tight
```

and displaying images

```
>> subplot(2,3,4)
```

```
>> imagesc(abs(x));axis off;
```

```
>> subplot(2,3,5)
```

```
>> imagesc(abs(xs));axis off;
```

```
>> subplot(2,3,6)
```

```
>> imagesc(abs(xRetrieved)); axis off;
```

5.4. Reproducing the experiments of the paper. All the numerical experiments of the paper can be reproduced using the Matlab scripts included in the toolbox directory `Experiments`.

- `phaseTransition_OSF1.m` (evolution of MSE with number of filters, with no oversampling of the Fourier transform, Figures 6, 7)
- `phaseTransition_OSF2.m` (evolution of MSE with number of filters, with oversampling of the Fourier transform Figures 8, 9)
- `filterResTransition.m` (evolution of MSE with filter resolution, figures 12, 13, 10, 11).
- `noiseTransition.m` (evolution of MSE with noise, Figures 14, 15)
- `testNoiseNbIllums.m` (test noise vs number of filters, Figures 3 and 4, and table 3)
- `testSeeds.m` (test different seeds to generate filters, Figure 5)

ACKNOWLEDGMENTS

AA and FF would like to acknowledge support from a starting grant from the European Research Council (project SIPA).

REFERENCES

- Heinz H Bauschke, Patrick L Combettes, and D Russell Luke. Phase retrieval, error reduction algorithm, and fienu variants: a view from convex optimization. *JOSA A*, 19(7):1334–1345, 2002.
- Stephen R Becker, Emmanuel J Candès, and Michael C Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218, 2011.
- Christian Berg, Jens Peter Reus Christensen, and Paul Ressel. *Harmonic analysis on semigroups : theory of positive definite and related functions*, volume 100 of *Graduate texts in mathematics*. Springer-Verlag, New York, 1984.
- H.M. Berman, T. Battistuz, TN Bhat, W.F. Bluhm, P.E. Bourne, K. Burkhardt, Z. Feng, G.L. Gilliland, L. Iype, S. Jain, et al. The protein data bank. *Acta Crystallographica Section D: Biological Crystallography*, 58(6):899–907, 2002.
- O. Bunk, A. Diaz, F. Pfeiffer, C. David, B. Schmitt, D.K. Satapathy, and JF Veen. Diffractive imaging for periodic samples: retrieving one-dimensional concentration profiles across microfluidic channels. *Acta Crystallographica Section A: Foundations of Crystallography*, 63(4):306–314, 2007.
- E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift : exact and stable signal recovery from magnitude measurements via convex programming. *To appear in Communications in Pure and Applied Mathematics*, 66(8): 1241–1274, 2013a.
- E.J. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *Arxiv preprint arXiv:1109.0573*, 2011.
- Emmanuel J Candès, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval from coded diffraction patterns. *preprint*, 2013b.
- A. Chai, M. Moscoso, and G. Papanicolaou. Array imaging using intensity-only measurements. *Inverse Problems*, 27:015005, 2011.
- Martin Dierolf, Andreas Menzel, Pierre Thibault, Philipp Schneider, Cameron M Kewish, Roger Wepf, Oliver Bunk, and Franz Pfeiffer. Ptychographic x-ray computed tomography at the nanoscale. *Nature*, 467(7314):436–439, 2010.
- J.R. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982.
- R. Gerchberg and W. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- Joseph Goodman. Introduction to fourier optics. 2008.
- D. Griffin and J. Lim. Signal estimation from modified short-time fourier transform. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 32(2):236–243, 1984.
- R.W. Harrison. Phase problem in crystallography. *JOSA A*, 10(5):1046–1055, 1993.
- C. Helmberg, F. Rendl, R. J. Vanderbei, and H. Wolkowicz. An interior–point method for semidefinite programming. *SIAM Journal on Optimization*, 6:342–361, 1996.
- I Johnson, K Jefimovs, O Bunk, C David, M Dierolf, J Gray, D Renker, and F Pfeiffer. Coherent diffractive imaging using phase front modifications. *Physical review letters*, 100(15):155503, 2008.
- L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- Filipe Maia. Spsim. 2013.
- J. Miao, T. Ishikawa, Q. Shen, and T. Earnest. Extending x-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes. *Annu. Rev. Phys. Chem.*, 59:387–410, 2008.
- N.Z. Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25:1–11, 1987.

K. C. Toh, M. J. Todd, and R. H. Tutuncu. SDPT3 – a MATLAB software package for semidefinite programming. *Optimization Methods and Software*, 11:545–581, 1999.

I. Waldspurger, A. d’Aspremont, and S. Mallat. Phase recovery, maxcut and complex semidefinite programming. *ArXiv: 1206.0102*, 2012.

C.M.A.P., ÉCOLE POLYTECHNIQUE, UMR CNRS 7641
E-mail address: fajwel.fogel@cmap.polytechnique.fr

D.I., ÉCOLE NORMALE SUPÉRIEURE, PARIS.
E-mail address: waldspur@clipper.ens.fr

CNRS & D.I., UMR 8548,
ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: aspremon@ens.fr

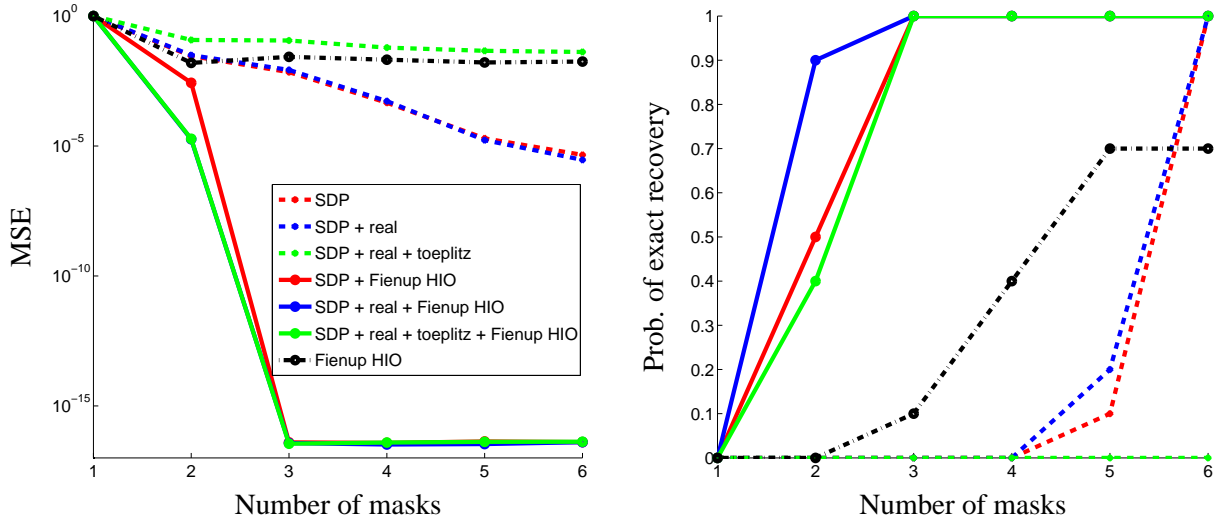


FIGURE 6. 16×16 caffeine image. No oversampling. *Left*: MSE (relative to $\|b\|$) vs. number of random masks. *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 7. 16×16 cocaine image. No oversampling. *Left*: MSE (relative to $\|b\|$) vs. number of random masks. *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.

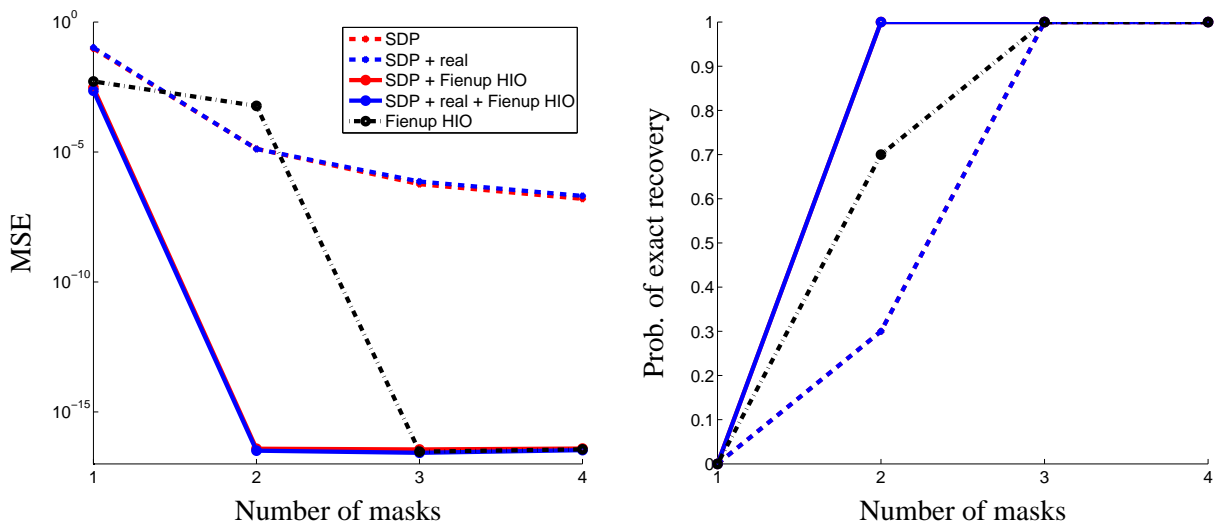


FIGURE 8. 16×16 caffeine image. 2x oversampling. *Left*: MSE vs. number of random masks. *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 9. 16×16 cocaine image. 2x oversampling. *Left*: MSE vs. number of random masks. *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 10. 16×16 caffeine image. Mask resolution (1x1 to 8x8 pixels). *Left*: MSE vs. mask resolution. (2x oversampling, no noise, 3 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 11. 16×16 cocaine image. Mask resolution (1x1 to 8x8 pixels). *Left*: MSE vs. mask resolution. (2x oversampling, no noise, 3 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 12. 16×16 caffeine image. Mask resolution (1x1 to 8x8 pixels). *Left*: MSE vs. mask resolution. (2x oversampling, no noise, 2 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.

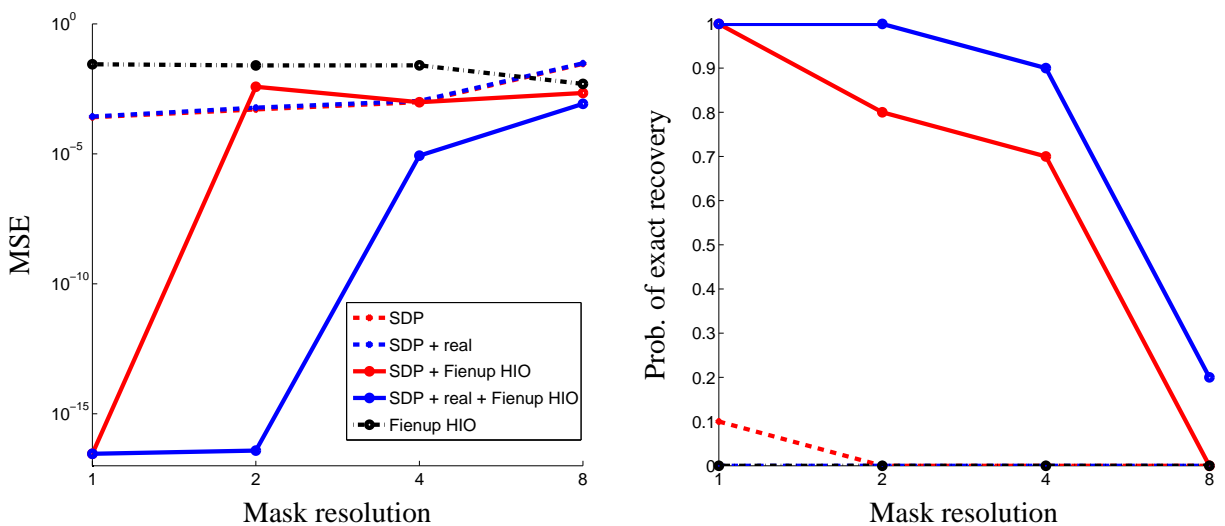


FIGURE 13. 16×16 cocaine image. Mask resolution (1x1 to 8x8 pixels). *Left*: MSE vs. mask resolution. (2x oversampling, no noise, 2 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.

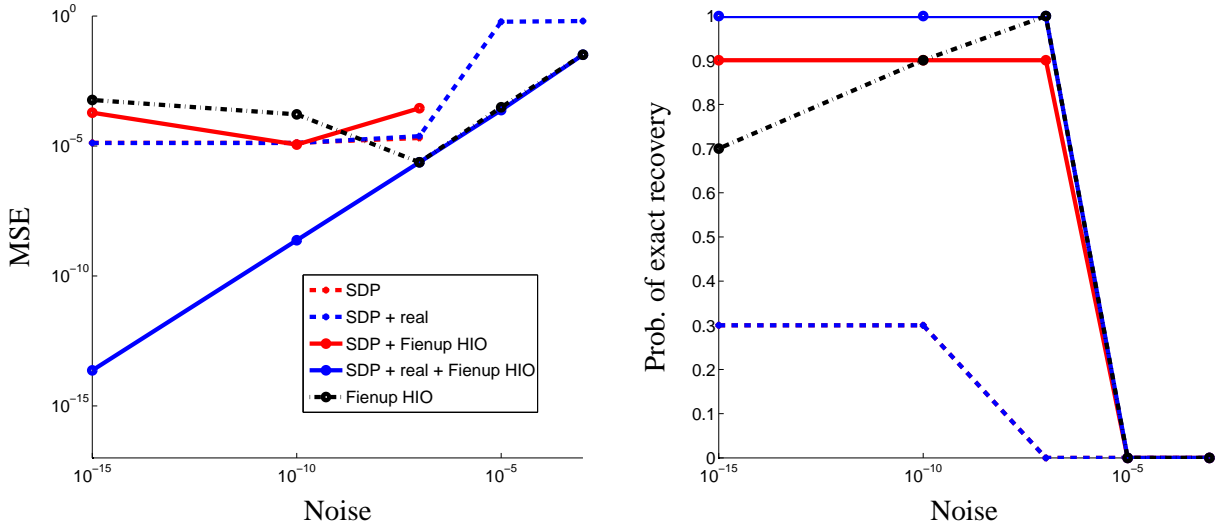


FIGURE 14. 16×16 caffeine image. Noise. *Left*: MSE vs. noise level α (2x oversampling, 2 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.




FIGURE 15. 16×16 cocaine image. Noise. *Left*: MSE vs. noise level α (2x oversampling, 2 masks). *Right*: Probability of recovering molecular density ($MSE < 10^{-4}$) vs. number of random masks.